

REMARKS/ARGUMENTS

The present Amendment is in response to the Office Action having a mailing date of November 23, 2005. Claims 1-16 are pending in the present Application. Applicant has amended claims 4-11. Consequently, claims 1-16 remain pending in the present Application.

Applicant has amended claims 1 and 15 to remove the phrase "capable of". Applicant has amended claim 5 to recite memory "usage" and LAN utilization in lieu of memory "using" and "LAN", respectively. Support for the amendment can be found in the specification, page 10, lines 2-11. Applicant has amended claim 4 to recite examples of future bottlenecks: latent bottlenecks and a cluster-level bottleneck. Support for the amendment can be found in the specification, page 10, line 12-page 11, line 15. Applicant has also amended claims 6-11 to more positively recite the cluster-level remedy. Accordingly, Applicant respectfully submits that no new matter is added.

In the above-identified Office Action, the Examiner rejected claim 14 under 35 U.S.C. § 101, as being directed toward non-statutory subject matter. In so doing, the Examiner stated that claim 14 recites "a computer readable medium. . . Applicant has defined computer readable media to further encompass a non-statutory 'computer readable signal which . . . may be transmitted over a network.' . . . Signals are not statutory subject matter."

Applicant respectfully disagrees with the Examiner's rejection. As stated in the specification:

Software written according to the present invention is to be stored in some form of computer-readable medium, such as memory, CD-ROM or transmitted over a network, and executed by a processor. Consequently, a computer-readable medium is intended to include a computer readable signal which, for example, may be transmitted over a network.

Specification, page 19, lines 4-8. Thus, as described in the specification, the signal is not considered to be separate from the medium, such as a network, over which the signal is transmitted.

Applicant notes that there are a variety of medium (for example including but not limited to a network, optical medium, electromagnetic medium, and infrared medium) over which such a signal may be transmitted. Notwithstanding the mention of a signal in the specification, the recitation of a computer readable medium in claim 14 is directed to the medium and, therefore, to statutory subject matter. Accordingly, Applicant respectfully submits that claim 14 is allowable under 35 U.S.C. § 101.

In the above-identified Office Action, the Examiner rejected claims 3, 4, 5, and 10 under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. In so doing, the Examiner cited the term “latent bottleneck” in claim 3, “future bottleneck” in claim 4, “portion of the workload” in claim 10, and “memory using” in claim 5.

Applicant respectfully disagrees with the Examiner’s rejection of claim 3. Claim 3 does recite a “latent bottleneck”. However, the specification defines a latent bottleneck as one which will occur in a cluster when another, current bottleneck is cleared. Specification, page 6, lines 13-20 and page 10, lines 13-23. Consequently, Applicant respectfully submits that the terms in claim 3 are clear and definite.

Applicant respectfully traverses the Examiner’s rejection of claim 4. Claim 4 recites a “future bottleneck.” As currently presented, claim 4 also recites examples of future bottlenecks. The specification discusses forecasting future bottlenecks (i.e. a bottleneck which does not currently exist, but which may exist in the future) in the context of latent bottlenecks (which do not currently exist but may come into existence once another bottleneck is cleared) and impending cluster level bottlenecks (in which failure of a node in a cluster may cause a bottleneck). Specification, page 10,

line 12-page 11, line 15 and Fig. 3, step 204. Consequently, Applicant respectfully submits that the terms in claim 4 are clear and definite.

Applicant respectfully traverses the Examiner's rejection of claim 5. Applicant has corrected claim 5 to recite "memory usage." Accordingly, Applicant respectfully submits that claim 5 is clear and definite.

Applicant respectfully disagrees with the Examiner's rejection of claim 10. Claim 10 recites that a node of the plurality of nodes carries a workload and has a bottleneck. Claim 10 also recites that a companion node in the cluster is "capable of supporting a portion of the workload." Stated differently, the companion node can support part of the workload of the bottlenecked node. Thus, claim 10 recites that the cluster remedy is capable of including a notification that the portion of the workload can be moved to the companion node. Stated differently, one cluster level remedy is a redistribution of the workload such that other node(s) carry a part of the workload of the bottlenecked node. Such a situation is described in the specification, page 11, line 18-page 12, line 14. Furthermore, the specific amount (e.g. a percentage of the workload carried by the bottlenecked node) to which the "portion" corresponds may depend upon the particularities of the implementation, such as the workload carried by each node and the number of nodes in the cluster. Furthermore, Applicant notes that the words "portion of the workload" have plain meanings and no further explanation is required. According to the MPEP, "[w]here elements ... and processes, which are conventional and generally widely known in the field to which the invention pertains, form a part of the invention described and their exact nature or type is not necessary for an understanding and use of the invention by a person skilled in the art, they should not be described in detail." MPEP 601(g). Consequently, Applicant respectfully submits that the term "portion of the workload" is well understood by one of ordinary

skill in the art and need not be further specified the in the claim in order to apprise one of ordinary skill in the art of the scope of the invention.

The Examiner also rejected claims 1-3, 5-7, 10, and 12-16 under 35 U.S.C. § 102 as being anticipated by U.S. Patent No. 6,434,626 (Prakash). The Examiner also rejected claims 4, 8-9, and 11 under Prakash in view of U.S. Patent No. 5,923,645 (Okuda). In the above-identified Office Action, the Examiner also rejected claim 1 over claim 1 of U.S. Patent No. 6,434,613 (Bertram).

The Examiner rejected claims 1-3, 5-7, 10, and 12-16 under 35 U.S.C. § 102 as being anticipated by Prakash. In so doing, the Examiner cited col. 3, lines 2-63 and col. 4, lines 21-33 and 47-65.

Applicant respectfully disagrees with the rejection. Claim 1 recites a method for providing performance analysis on a system including a cluster, which includes a plurality of nodes. The method includes dynamically obtaining data relating to monitors for the plurality of nodes in the cluster by sampling the plurality of nodes and dynamically analyzing the data to determine whether performance of the cluster can be improved. The method further includes providing at least one remedy to improve performance of the cluster if the performance of the cluster can be improved, the at least one remedy including a cluster level remedy. Claims 14 and 15 recite an analogous computer readable medium and an analogous system. Claim 13 recites an analogous method that further specifically accounts for computer systems not in the cluster.

Applicant respectfully disagrees that Prakash teaches the methods, system, and computer readable medium recited in claims 1 and 13-15. In particular, Prakash fails to teach or suggest providing a cluster level remedy if the results of dynamic analysis indicate that the performance of the cluster can be improved. The cited portions of Prakash describe the general benefits of using a clusters and summary of the invention described in Prakash. Prakash, col. 3, line 1-col.

5, line 9. In describing adding data paths, increasing point-to-point data paths, establishing peer-to-peer communications, and other benefits of clusters, Prakash is contrasting clusters to other conventional networks. In particular, Prakash states:

Another solution for relieving network bottlenecks and achieving scalability [separate from the previously discussed I/O architecture] is to provide a clustered network environment wherein a variety of components like servers, disk drives, tape drives, etc., are integrated into a system-wide architecture such as a System Area Network (SAN). SAN architectures, for example, a fabric network, provide a low latency interconnect between servers and devices and can be configured for I.sub.2.O compliance. SAN architecture is based on message passing between servers and devices. SAN technology employs the server processor to process data transfer requests between network elements and then allow data transfers to occur under control of dedicated hardware thus reducing server processor overhead to a minimum. In a SAN architecture, a network transport layer may be implemented on a dedicated hardware platform, typically an I/O processor (IOP), which allows a processor to be connected to a scalable switching fabric. A SAN server can then be expanded to add data paths which effectively increase the overall bandwidth of the switching fabric by increasing the number of point-to-point datapaths which can be used to carry data between nodes. Thus, large numbers of nodes which may be clients, other servers, or other network devices such as disk towers, and the like may be controlled by a server. Further, to off-load the processing of data transfers from the server processor, peer-to-peer communications may be set up between devices and the transfers may proceed without further server intervention.

Prakash, col. 2, line 62-col. 3, line 22 (emphasis added). Thus, the possibility of setting up different types of communication, off-loading of processing and other benefits of cluster technology are generally discussed. Prakash goes on to state that performance is monitored in such a network by indicating whether devices are in a healthy or unhealthy state, and action such as issuing error messages taken in response to indications that a device is "unhealthy." Prakash, col. 3, lines 22-34. Prakash then describes the problems associated with monitoring performance of devices in such a network without indicating how to address such problems or whether remedies, including cluster level remedies, are provided. Prakash, col. 3, line 35-col. 4, line 12.

The Summary of the Invention, a portion of which was also cited by the Examiner, describes a method of addressing such problems. In particular, Prakash describes sending inquiries and receiving messages indicating whether a device is “healthy” and error free or “unhealthy” and thus malfunctioning in some manner. Prakash, col. 4, lines 52-63. Prakash goes on to state that if an “unhealthy” message is received, the response is to establish further communication to request detailed information. Prakash, col. 4, line 63-col. 5, line 1.

Thus, Prakash does describe various benefits of cluster architecture as well as mechanisms for monitoring whether devices in the network are “healthy” or “unhealthy”. However, the cited portions of Prakash do not indicate that any remedy, much less a cluster level remedy, is provided in response to performance issues. Consequently, Prakash fails to teach or suggest the method, system, and computer-readable medium recited in claims 1 and 13-15.

Claims 2-3, 5-7, 10, and 12 depend upon independent claim 1. Claim 16 depends upon independent claim 15. Consequently, the arguments herein apply with full force to claims 2-3, 5-7, 10, 12, and 16. Accordingly, Applicant respectfully submits that claims 2-3, 5-7, 10, 12, and 16 are allowable over the cited references.

Claims 3, 5-7, 10, and 12 are also separately allowable over the cited references. Claim 3 recites determining whether a latent bottleneck exists. As discussed above, a latent bottleneck is one which may exist once a current bottleneck is cleared. Although, as the Examiner has indicated, Prakash indicates whether a device is “healthy” or “unhealthy”, Applicant has found no indication in Prakash that the existence of a latent bottleneck is specifically included in determining the health of a device. Consequently, Prakash fails to teach or suggest the method recited in claim 3.

Claim 5 specifies that the monitors for which data are dynamically obtained and analyzed include disk utilization, CPU utilization, memory usage and LAN utilization. The cited portion of Prakash merely indicates that a “healthy” device is one which is error-free, while an “unhealthy” device is malfunctioning in some manner. Prakash, col. 4, lines 58-62. Applicant respectfully submits that an indication of being “unhealthy” or malfunctioning could include any number of issues which may not be related to the recited monitors. Consequently, without more, Prakash fails to teach or suggest the method recited in claim 5.

Claims 6-7 and 10 recite various cluster-level remedies that may be provided if performance of the cluster performance can be improved. Applicant notes that, as discussed above, the portions of Prakash cited by the Examiner generally discuss the benefits of cluster architecture, not remedies that might be provided to enhance performance in response to dynamic analysis of data related to monitors in the cluster. Consequently, Prakash fails to teach or suggest the methods recited in claims 6-7 and 10.

Claims 12 recites a method that also includes obtaining information obtaining information relating to the cluster. This information specifically includes an indication of whether each of the plurality of nodes is a passive node, a maximum number of nodes in the cluster and a type of LAN adapter used for interconnecting the plurality of nodes. The portion of Prakash cited by the Examiner indicates that the responses to status messages include only two possibilities: healthy and unhealthy. Prakash, col. 4, lines 52-67. Consequently, the specifics relating to nodes of the cluster are not part of the information passed in response to status inquiries. Prakash thus fails to teach or suggest the method recited in claim 12. Accordingly, for the above-identified reasons, claims 5-7, 10, and 12 are separately allowable over the cited references.

The Examiner also rejected claims 4, 8-9, and 11 under Prakash in view of U.S. Patent No. 5,923,645 (Okuda).

Applicant respectfully disagrees with the Examiner's rejection. Claims 4, 8-9, and 11 depend upon independent claim 1. Consequently, the arguments herein with respect to Prakash apply with full force to claims 4, 8-9, and 11. Thus, Prakash fails to teach or suggest dynamically obtaining and analyzing data for a plurality of monitors in conjunction with providing a remedy, at least one of which may be a cluster level remedy.

The cited portion of Okuda does mention predicting a bottleneck. However, the cited portion of Okuda is devoid of mention of cluster level remedies. Consequently, even if the ability to predict certain bottlenecks is added to the teachings of Prakash, the combination would still fail to teach or suggest the recited cluster level remedies. Furthermore, Applicant respectfully submits that any conclusion that Prakash in view of Okuda teaches or suggests the methods recited in claims 4, 8-9 and 11 involves improper hindsight. Applicant notes that one "cannot use hindsight reconstruction to pick and choose among isolated disclosures in the prior art to deprecate the claimed invention." In re Fine, 5 USPQ2d 1596, 1600 (Fed. Cir. 1988). See also In re Fritch, 23 USPQ2d 1780, 1783 (Fed. Cir. 1992). Consequently, claims 4, 8-9, and 11 are allowable over the cited references.

Furthermore, claims 8-9 and 11 are separately allowable over the cited references. Claims 8-9 and 11 recite various cluster level remedies such as warning that if a node may become bottlenecked if another node in the cluster fails, indicating that a companion node may be a source of a bottleneck if another node is bottlenecked, and a notification that a node will become bottlenecked if another node fails. Applicant has found no mention in the cited portions of Prakash or Okuda of such specific cluster level remedies. Consequently, any combination of

the cited portions of Prakash and Okuda would also fail to teach or suggest such remedies. Accordingly, Applicant respectfully submits that claims 8-9 and 11 are separately allowable over the cited references.

In the above-identified Office Action, the Examiner also rejected claim 1 over claim 1 of Bertram based on the judicially created doctrine of double patenting.

Applicant respectfully disagrees with the Examiner's rejection. Applicant has found no mention in claim 1 of Bertram providing cluster level remedies. Consequently, although Bertram functions well for its intended purpose, Applicant respectfully submits that claim 1 of the present application is patentably distinct from claim 1 of Bertram. Accordingly, Applicant respectfully submits that claim 1 is allowable over the cited references.

In the above-identified Office Action, the Examiner also rejected claim 1 over claim 1 of U.S. Patent No. 6,470,464 (Bertram II) on the grounds of statutory double patenting.

Applicant respectfully disagrees with the Examiner's rejection. Statutory double patenting is same invention type double patenting. In this context same invention means identical subject matter. *Miller v. Eagle Mfg. Co.*, 151 U.S. 186 (1984); *In re Vobel*, 422 F.2d 438, 164 USPQ 619 (CCPA 1970); and *In re Ockert*, 245 F.2d 467, 114 USPQ 330. Thus, "[w]here the claims of an application are not the 'same' as those of a first patent, but the grant of a patent would unjustly extend the right granted by the first patent, a double patenting rejection under nonstatutory grounds is proper." MPEP 804.II (page 800-19). In addition:

[d]omination and double patenting should not be confused. They are two separate issues. One patent or application "dominates" a second patent or application when the first patent or application has a broad or generic claim which fully encompasses or reads on an invention defined in a narrower or more specific claim in another patent or application. Domination by itself, i.e., in the absence of

statutory or nonstatutory double patenting grounds cannot support a double patenting rejection.

MPEP 804.II (page 800-19).

Claim I recites:

1. A method for providing performance analysis on a system including a cluster, the cluster including a plurality of nodes, the method comprising the steps of:

 dynamically obtaining data for the plurality of nodes in the cluster by sampling the plurality of nodes, the data relating to a plurality of monitors for the node,

 dynamically analyzing the data to determine whether performance of the cluster can be improved;

 providing at least one remedy to improve performance of the cluster if the performance of the cluster can be improved, the at least one remedy including a cluster level remedy.

In contrast, claim 1 of Bertram II recites:

1. A method of analyzing performance of a computer system and for providing recommendations for changes in the system to improve its performance, the steps of the method comprising:

 receiving and storing indicators of the performance of components of the computer system at periodic times;

 applying stored rules to the stored indicators to analyze the performance of the computer system; project future performance of the computer system based upon past indicators;

 detecting a projected undesirable condition of the computer system based upon the future performance projections based on past indicators;

 in response to the detecting of an undesirable condition of the computer system, providing a recommendation to alleviate the performance of computer system by ameliorating the projected undesirable conditions of the computer system; and

 analyzing the undesirable conditions to determine the most severe undesirable condition and reporting it first.

Thus, Bertram II recites detecting a projected undesirable condition based on future performance projections of past indicators and in response, providing a recommendation. There is no indication that such a recommendation must include a cluster level remedy. In particular, a cluster level

remedy is defined in the specification as one which is capable of being performed for a cluster, but not for a system having only a single node. Specification, page 11, line 15-page 12, line 18. For example, cluster level remedies may include moving resources between nodes, adding nodes, or warning that a particular node may fail so that the user can make changes to the cluster and the node's workload need not be absorbed by remaining nodes. Ibid. Thus, for example, adding more memory to a particular system may not be a cluster level remedy, but transferring jobs between systems may be.

Bertram II describes recommendations that are not cluster level remedies. For example, in Bertram II, a disk bottleneck, a CPU bottleneck, and a memory bottleneck are described. The recommendations for these bottlenecks include adding more disks, upgrading the CPU, and adding more memory. Applicant respectfully submits that such remedies are not "cluster-level" remedies. Thus, if a particular embodiment provided these recommendations of Bertram II, but not cluster-level remedies, claim 1 of Bertram II, but not claim 1 of the present application, may be infringed. Consequently, the inventions as defined by claim 1 of the present application and claim 1 of Bertram are not the same. Claim 1 of the present application thus does not run afoul of statutory double patenting of claim 1 of Bertram II. Accordingly, Applicant respectfully submits that claim 1 is allowable over Bertram II.

Applicant's attorney believes that this application is in condition for allowance. Should any unresolved issues remain, Examiner is invited to call Applicant's attorney at the telephone number indicated below.

Respectfully submitted,
SAWYER LAW GROUP LLP

May 30, 2006
Date

Janyce R. Mitchell/ Reg. No. 40,095
Janyce R. Mitchell
Attorney for Applicant(s)
(650) 493-4540